

CLAIMS

What is claimed is:

1. An interconnection network, comprising:
 - 5 routers interconnected by links;
packets traversing one or more hops across links from packet sources to packet destinations;
subtrees of the interconnected routers forming source based virtual networks for each of the packet sources;
 - 10 subtrees of the interconnected routers forming egress based virtual networks for each of the packet destinations; and
a packet traversing a source based virtual network of a packet source, the packet transitioning into an egress based virtual network of a packet destination, the packet traversing the egress based virtual network
 - 15 to the packet destination.
2. The interconnection network of claim 1, wherein the routers comprise dynamically assignable buffer resources for implementing source based virtual networks and egress based virtual networks.
3. The interconnection network of claim 1, wherein each of the routers comprise
 - 20 memory implementing sets of queues for source based virtual networks and egress based virtual networks, the sets of queues for managing the transmission of packets over the virtual networks.
4. The interconnection network of claim 1, wherein:
 - 25 a source based virtual network fans out from a common packet source over a subtree of interconnected routers through source based

tunnels, each of the source based tunnels extending less than the diameter of the interconnection network; and

an egress based virtual network fans in to a common packet destination over a subtree of interconnected routers through egress based tunnels, each of the egress based tunnels extending less than the diameter of the interconnection network.

- 5
- 10
- 15
- 20
- 25
- 09066335-052504
5. The interconnection network of claim 4, wherein a packet is communicated from a packet source to a packet destination by traversing a source based tunnel associated with the packet source, transitioning into an egress based tunnel associated with the packet destination, and traversing the egress based tunnel to the packet destination.
 6. The interconnection network of claim 4, wherein:
 - each of the source based tunnels is a logical path over a set of interconnected routers, each of the source based tunnels is associated with a common packet source, each of the source based tunnels extending zero or more hops from the packet source to a binding node; and
 - each of the egress based tunnels is a logical path over a set of interconnected routers, each of the egress based tunnels is associated with a common packet destination, each of the egress based tunnels extending zero or more hops from a binding node to a common packet destination.
 7. The interconnection network of claim 4, wherein the routers comprise dynamically assignable buffer resources for implementing source based tunnels and egress based tunnels.

8. The interconnection network of claim 4, wherein each of the routers comprise memory implementing sets of queues for source based tunnels and egress based tunnels, the sets of queues for managing the transmission of packets over the tunnels.
9. The interconnection network of claim 1, wherein the packet sources and the packet destinations are ports.
10. The interconnection network of claim 1, wherein the packet sources and the packet destinations are data channels.
11. The interconnection network of claim 10, wherein the data channels are lanes.
12. The interconnection network of claim 1, wherein the packet sources and the packet destinations are groups of multiple data channels.
13. The interconnection network of claim 1 is a fabric of an Internet router.
14. The interconnection network of claim 13 is a vertex symmetric, direct fabric.
15. The interconnection network of claim 13 is a Gamma graph fabric.
16. The interconnection network of claim 1 is a fabric of a multi-application switch router.
17. The interconnection network of claim 16 is a vertex symmetric, direct fabric.
18. The interconnection network of claim 16 is a Gamma graph fabric.

19. An interconnection network, comprising:
- 5 routers interconnected by links;
 - packets traversing one or more hops across links from packet sources to packet destinations;
 - 10 a plurality of source based tunnels fanning out from each packet source over a subtree of routers, each source based tunnel comprises one or more source based tunnel segments, each source based tunnel segment of a source based tunnel is a logical hop across a different link; and
 - a plurality of egress based tunnels fanning in to each packet destination over a subtree of routers, each egress based tunnel comprises one or more egress based tunnel segments, each egress based tunnel segment of an egress based tunnel is a logical hop across a different link.
20. The interconnection network of claim 19, wherein a source based tunnel segment and an egress based tunnel segment is dynamically assigned buffer resources at an adjacent fabric router.
21. The interconnection network of claim 19, wherein each router comprises memory implementing sets of queues for source based tunnel segments and egress based tunnel segments, the set of queues for managing the transmission of packets over the tunnel segments.
22. The interconnection network of claim 19, wherein:
- each router comprises memory implementing a queue for each egress based tunnel segment associated with a packet destination that is reachable from the router in a number of hops less than the diameter of the interconnection network;
 - 25 each router comprises memory implementing a queue for each source based tunnel segment associated with a packet source and a path

FOIA b 7 - DUE TO THE PRESENCE OF

to a binding node that is reachable from the router in a number of hops less than the diameter of the interconnection network; and

the total number of queues for source based tunnel segments and egress based tunnel segments at each router being less than the total number of packet destinations of the interconnection network.

23. The interconnection network of claim 19, wherein packets traverse through source based tunnel segments and egress based tunnel segments one packet at a time.

24. The interconnection network of claim 19, wherein a packet traversing a tunnel segment of a congested tunnel is bypassed by another packet traversing a tunnel segment of a non-congested tunnel even if the tunnel segment of the congested tunnel and the tunnel segment of the non-congested tunnel correspond to the same path segment.

25. The interconnection network of claim 19, wherein the routers comprise:
a switch coupled to at least one traffic manager, each traffic manager coupled to at least one packet source and at least one packet destination, each switch being coupled to one or more routers by links;

the traffic manager injecting packets from the at least one packet source into one of the plurality of source based tunnels and ejecting packets from one of the plurality of egress based tunnels to the at least one packet destination, the switch forwarding packets to switches of next hop routers through source based tunnel segments and egress based tunnel segments, the switch forwarding packets destined for the at least one packet destination to the traffic manager.

26. The interconnection network of claim 25 wherein the traffic manager is coupled to the at least one packet source and the at least one packet destination by an Infiniband™ bus.
27. The interconnection network of claim 25 wherein the traffic manager is coupled to the at least one packet source and the at least one packet destination by a CSIX™ bus.
28. An interconnection network, comprising:
- router interconnected by links;
 - packets traversing one or more fabric hops across links from packet sources to packet destinations;
 - each router comprising a switch coupled to at least one traffic manager, each traffic manager coupled to at least packet source and at least one packet destination, the switch communicating with the traffic manager and one or more routers;
 - each traffic manager injecting packets from the at least one packet source into source based tunnel segments, each of the source based tunnel segments associated with a packet source and a path to a binding node originating an egress based tunnel to at least one packet destination; and
 - the switch forwarding packets through source based tunnel segments, the switch forwarding packets through egress based tunnel segments, the egress based tunnel segments corresponding to packet destinations reachable in a number of hops less than the diameter of the interconnection network.

29. The interconnection network of claim 28, wherein the traffic manager prefixes packets from packet destinations with a header, the header specifying a switch egress port for each hop along the path from a packet source to a packet destination.
- 5 30. The interconnection network of claim 28, wherein the traffic manager segments each packet into fixed size data blocks.
31. The interconnection network of claim 30, wherein the fixed size data blocks are flits, each flit comprising a portion of the packet and a flit header.
32. The interconnection network of claim 28 wherein:
- 10 the traffic manager comprises an egress controller, the egress controller manages a plurality of queues corresponding to source based tunnel segments for packet transport to the switch, the queues referencing packets waiting to transit the source based tunnel segments; and
- the switch comprises an egress controller for each switch egress
- 15 port, each egress controller manages a plurality of queues corresponding source based tunnel segments or egress based tunnel segments for packet transport across a link interconnecting adjacent routers, the queues referencing packets waiting to transit the source based tunnel segments and the egress based tunnel segments.
- 20 33. The interconnection network of claim 32, wherein each queue is dedicated to a source based tunnel or an egress based tunnel.
34. The interconnection network of claim 32, wherein:
- the plurality of source based tunnel segment queues are shared among source based tunnels; and

FOIA b 7 - DATED 05/23/2011 BY SP5/BJP

the plurality of egress based tunnel segment queues are shared among egress based tunnels.

35. The interconnection network of claim 32, wherein:

5 an egress controller injects a packet into a tunnel segment when the tunnel segment is not busy and buffer resources are available to the tunnel segment.

36. The interconnection network of claim 32 wherein each link comprising a set of lanes, each lane corresponding to two or more buffers associated with a switch ingress port of a next hop router, each buffer capable of storing at least a portion of a packet.

10

37. The interconnection network of claim 36 wherein the set of lanes is partitioned into lane groups, each lane group corresponding to a hop distance, an egress controller of the switch allocating a non-busy lane in a lane group to a tunnel segment if the destination of a data packet traversing the tunnel segment is at a hop distance away from the egress controller that is equal to the hop distance corresponding to the lane group.

15

38. The interconnection network of claim 36, wherein the egress controller of the traffic manager and the egress controllers of the switch dynamically allocate an available lane to each tunnel segment waiting to transport packets, an available lane being assignable to any tunnel segment for transporting packets.

20

39. The interconnection network of claim 36 wherein a lane is released from a tunnel segment when the data blocks of a data packet have been forwarded across the next hop fabric router.

"05250" SEC99999

40. A method for preventing tree saturation and deadlock in an interconnection network, comprising:
- interconnecting routers by links, each router coupled to one or more packet sources and one or more packet destinations;
 - 5 forming source based virtual networks over subtrees of interconnected routers for each of the packet sources;
 - forming egress based virtual networks over subtrees of interconnected routers for each of the packet destinations; and
 - 10 communicating a packet from any packet source to any packet destination by traversing a source based virtual network of a packet source, transitioning into an egress based virtual network of a packet destination, and traversing the egress based virtual network to the packet destination.
41. The method of claim 40, further comprising:
- 15 dynamically assigning buffer resources of a router to a source based or egress based virtual network for buffering a packet that is traversing across the router over the source based or egress based virtual network.
42. The method of claim 40, further comprising:
- 20 implementing sets of queues within memory of each of the routers for source based virtual networks and egress based virtual networks; and
 - managing transmission of packets over the source based and egress based virtual networks with each queue referencing zero or more
 - 25 packets waiting to traverse a virtual network.

F05250"SEE9990

43. The method of claim 40, forming a source based virtual network and an egress based virtual network further comprise:

associating a set of source based tunnels with a common packet source that fan out over a subtree of interconnected routers, each of the source based tunnels extending less than the diameter of the interconnection network; and

associating a set of egress based tunnels with a common packet destination that fan in over a subtree of interconnected routers to the packet destination, each of the egress based tunnels extending less than the diameter of the interconnection network.

44. The method of claim 43, wherein communicating a packet from a packet source to a packet destination, further comprises:

traversing a source based tunnel associated with the packet source;

transitioning into an egress based tunnel associated with the packet destination; and

traversing the egress based tunnel to the packet destination.

45. The method of claim 43, further comprising:

dynamically assigning buffer resources of a router to a source based or egress based tunnel for buffering a packet traversing across the router over the source based or egress based tunnel.

46. The method of claim 43, further comprising:

implementing sets of queues within memory of each of the routers for source based tunnels and egress based tunnels; and

FOIA b 5 - DECLASSIFIED

managing transmission of packets over the source based and egress based tunnels with each queue referencing zero or more packets waiting to traverse a tunnel.

47. A method for preventing tree saturation and deadlock in an interconnection network, comprising:
- 5 interconnecting routers by links, each router coupled to one or more packet sources and one or more packet destinations;
- 10 for each router, implementing a queue within memory for each egress based tunnel segment associated with a packet destination that is reachable from the router in a number of hops less than the diameter of the interconnection network, an egress based tunnel segment is a logical hop across a link associated with a packet destination;
- 15 for each router, implementing a queue within memory for each source based tunnel segment associated with a packet source and a path to a binding node that is reachable from the router in a number of hops less than the diameter of the interconnection network, a source based tunnel segment is a logical hop across a link associated with a packet source; and
- 20 the total number of queues for source based tunnel segments and egress based tunnel segments at each router being less than the total number of packet destinations of the interconnection network.
48. The method of claim 47, further comprising:
- 25 dynamically assigning buffer resources of a router to a source based or egress based tunnel segment for buffering a packet traversing across the router over the source based or egress based tunnel segment.

0986336-052501

49. The method of claim 47, further comprising:
implementing sets of queues within memory of each of the
routers for source based tunnel segments and egress based tunnel
segments; and
5 managing transmission of packets over the source based and
egress based tunnel segments with each queue referencing zero or more
packets waiting to traverse a tunnel segment.
50. The method of claim 47, further comprising:
forwarding packets through source based and egress based tunnel
10 segments one packet at a time; and
forwarding packets through tunnel segments of non-congested
tunnels bypassing packets waiting to traverse tunnel segments of
congested tunnels even if the tunnel segments of the congested tunnel
and the non-congested tunnel correspond to the same link.
- 15 51. The method of claim 48, wherein dynamic assignment of buffer resources for
packet buffering, further comprises:
partitioning each link into a set of lanes, each lane corresponding
to two or more buffers, each buffer capable of storing at least a portion of
a packet;
20 partitioning the set of lanes into lane groups, each lane group
corresponding to a hop distance;
allocating a non-busy lane to a tunnel segment from a lane group
associated with a hop distance that corresponds to the hop count of a
packet traversing the tunnel segment; and
25 forwarding the packet across the link storing the packet in the
buffers associated with the lane;

T05350" 5E99369

releasing the lane when the packet has been forwarded across the next hop.

52. An interconnection network, comprising:

routers interconnected by links;

5 packets traversing one or more hops across links from packet sources to packet destinations;

subtrees of the interconnected routers forming sets of source based tunnels associated with each of the packet sources;

10 subtrees of the interconnected routers forming sets of egress based tunnels associated with each of the packet destinations;

each router having a means for managing the transmission of packets through each source based tunnel and each egress based tunnel originating from or extending across the router; and

15 communicating a packet from any packet source to any packet destination by traversing a source based tunnel associated with the packet source, the packet transitioning into an egress based tunnel associated with the packet destination, the packet traversing the egress based virtual network to the packet destination.

20

2025-05-29 09:59:59